

1 METHODS AND APPARATUS FOR IMPLEMENTING LCAS (LINK CAPACITY
2 ADJUSTMENT SCHEME) SINKING WITH RATE BASED FLOW CONTROL

3
4 BACKGROUND OF THE INVENTION

5
6 1. Field of the Invention

7 The invention relates to telecommunications, the Synchronous
8 Optical Network (SONET) and the Synchronous Digital Hierarchy
9 (SDH). More particularly, the invention relates to methods and
10 apparatus for preventing the loss of data at an LCAS sink.

11
12 2. State of the Art

13 The Synchronous Optical Network (SONET) or the Synchronous
14 Digital Hierarchy (SDH), as it is known in Europe, is a common
15 telecommunications transport scheme which is designed to
16 accommodate both DS-1 (T1) and E1 traffic as well as multiples
17 (DS-3 and E-3) thereof. A DS-1 signal consists of up to twenty-
18 four time division multiplexed DS-0 signals plus an overhead bit.
19 Each DS-0 signal is a 64kb/s signal and is the smallest allocation
20 of bandwidth in the digital network, i.e. sufficient for a single
21 telephone connection. An E1 signal consists of up to thirty-two
22 time division multiplexed DS-0 signals with at least one of the
23 DS-0s carrying overhead information.

1 Developed in the early 1980s, SONET has a base (STS-1) rate
2 of 51.84 Mbit/sec in North America. The STS-1 signal can
3 accommodate 28 DS-1 signals or 21 E1 signals or a combination of
4 both. The basic STS-1 signal has a frame length of 125
5 microseconds (8,000 frames per second) and is organized as a frame
6 of 810 octets (9 rows by 90 byte-wide columns). It will be
7 appreciated that $8,000 \text{ frames} \times 810 \text{ octets per frame} \times 8 \text{ bits per}$
8 $\text{octet} = 51.84 \text{ Mbit/sec}$. The frame includes the synchronous
9 payload envelope (SPE) or virtual container (VC) as it is known in
10 Europe, as well as transport overhead. Transport overhead is
11 contained in the first three columns (27 bytes) and the SPE/VC
12 occupies the remaining 87 columns.

13

14 In Europe, the base (STM-1) rate is 155.520 Mbit/sec,
15 equivalent to the North American STS-3 rate ($3 \times 51.84 = 155.520$).
16 The STS-3 (STM-1) signals can accommodate 3 DS-3 signals or 63 E1
17 signals or 84 DS-1 signals, or a combination of them. The STS-12
18 signals are 622.080 Mbps and can accommodate 12 DS-3 signals, etc.
19 The STS-48 signals are 2,488.320 Mbps and can accommodate 48 DS-3
20 signals, etc. The highest defined STS signal, the STS-768, is
21 nearly 40 Gbps (gigabits per second). The abbreviation STS stands
22 for Synchronous Transport Signal and the abbreviation STM stands
23 for Synchronous Transport Module. STS-n signals are also referred

1 to as Optical Carrier (OC-n) signals when transported optically
2 rather than electrically.

3
4 To facilitate the transport of lower-rate digital signals,
5 the SONET standard uses sub-STs payload mappings, referred to as
6 Virtual Tributary (VT) structures. (The ITU calls these Tributary
7 Units or TUs.) This mapping divides the SPE (VC) frame into seven
8 equal-sized sub-frames or VT (TU) groups with twelve columns of
9 nine rows (108 bytes) in each. Four virtual tributary sizes are
10 defined as follows.

11
12 VT1.5 has a data transmission rate of 1.728 Mb/s and
13 accommodates a DS1 signal with overhead. The VT1.5 tributary
14 occupies three columns of nine rows, i.e. 27 bytes. Thus, each VT
15 Group can accommodate four VT1.5 tributaries.

16
17 VT2 has a data transmission rate of 2.304 Mb/s and
18 accommodates a CEPT-1 (E1) signal with overhead. The VT2
19 tributary occupies four columns of nine rows, i.e. 36 bytes.
20 Thus, each VT Group can accommodate three VT2 tributaries.

21
22 VT3 has a data transmission rate of 3.456 Mb/s) and
23 accommodates a DS1C (T2) signal with overhead. The VT3 tributary

1 occupies six columns of nine rows, i.e. 54 bytes. Thus, each VT
2 Group can accommodate two VT3 tributaries.

3
4 VT6 has a data transmission rate of 6.912 Mb/s and
5 accommodates a DS2 signal with overhead. The VT6 tributary
6 occupies twelve columns of nine rows, i.e. 108 bytes. Thus, each
7 VT Group can accommodate one VT6 tributary.

8
9 As those skilled in the art will appreciate, the original
10 SONET/SDH scheme as well as the VT mapping schemes were designed
11 to carry known and potentially foreseeable TDM signals. In the
12 early 1980s these TDM signals were essentially multiplexed
13 telephone lines, each having the (now considered) relatively small
14 bandwidth of 56-64kbps. At that time, there was no real standard
15 for data communication. There were many different schemes for
16 local area networking and the wide area network which eventually
17 became known as the Internet was based on a "56k backbone". Since
18 then, Ethernet has become the standard for local area networking.
19 Today Ethernet is available in four bandwidths: the original 10
20 Mbps system, 100 Mbps Fast Ethernet (IEEE 802.3u), 1,000 Mbps
21 Gigabit Ethernet (IEEE 802.3z/802.3ab), and 10 Gigabit Ethernet
22 (IEEE 802.3ae).

1 In recent years it has been recognized that SONET/SDH is the
2 most practical way to link high speed Ethernet networks over a
3 wide area. Unfortunately, the various Ethernet transmission rates
4 (10 Mbps, 100 Mbps, 1,000 Mbps, and 10,000 Mbps) do not map well
5 into the SONET/SDH frame. For example, the original 10 Mbps
6 Ethernet signal is too large for a VT-6 tributary but too small
7 for an entire STS-1. In other words, under the existing SONET/SDH
8 schemes, in order to transport a 10 Mbps Ethernet signal, an
9 entire STS-1 path must be used, thereby wasting a significant
10 amount of bandwidth. Similar results occur when attempting to map
11 the faster Ethernet signals into STS signals.

12
13 In order to provide a scheme for efficiently mapping Ethernet
14 signals (as well as other signals such as Fiber Channel and ESCON)
15 into a SONET/SDH frame, the Virtual Concatenation Protocol was
16 created and has been endorsed by the ITU as the G.707 standard.
17 Similar to inverse multiplexing, Virtual Concatenation combines
18 multiple links (members) into one Virtual Concatenation Group
19 (VCG), enabling the carrier to optimize the SDH/SONET links for
20 Ethernet traffic. For example, using virtual concatenation, five
21 VT-2 (2 Mbps) links can be combined to carry a 10 Mbps Ethernet
22 signal, resulting in full utilization of allotted bandwidth. Two
23 STS-1 (51 Mbps) links can be combined to carry a 100 Mbps Ethernet
24 signal, etc. Virtual Concatenation uses SONET/SDH overhead bytes

1 (four of the sixteen "H4" bytes) to indicate two numbers: the
2 multiframe indicator (MFI) and the sequence number (SQ).

3
4 The different tributaries that make up a VCG are allowed to
5 follow different paths through the network. This means that group
6 members that leave a network node at the same time and properly
7 interleaved will arrive at their destination at different times
8 with respect to each other. It is the responsibility of the
9 receiver to remove this differential delay (or skew) between group
10 members and restore their proper order before demapping the
11 payload contents.

12
13 Part of the Virtual Concatenation Protocol includes methods
14 for dynamically scaling the available bandwidth in a SONET/SDH
15 signal. These methods are known as the Link Capacity Adjustment
16 Scheme or LCAS. LCAS is a powerful network management tool
17 because customer bandwidth requirements change over time. One
18 simple example is a network user who, during business hours, needs
19 only enough bandwidth to support electronic mail and worldwide web
20 access. During non-working hours, however, the same network user
21 may wish to conduct relatively large data transfers from one
22 location to another to backup daily transactions, for example. It
23 would be desirable to alter the user's available bandwidth as
24 needed. LCAS provides a means to do this without disturbing other

1 traffic on the link. LCAS has been endorsed by the ITU as the
2 G.7042 standard which is hereby incorporated by reference herein
3 in its entirety.

4
5 LCAS creates a bidirectional communication scheme between a
6 transmitter and a receiver (the source and sink respectively) that
7 allows the size of a VCG to be modified "on the fly",
8 theoretically with no disruption of data. The source, through an
9 in-band communication channel, sends a message to the sink that
10 group members are about to be added or removed, then sends a
11 synchronization message so that the sink knows exactly when the
12 change will take effect. Removing group members in the presence
13 of differential delay creates some special problems for
14 maintaining data integrity.

15
16 A typical system 10 for recovering packetized (e.g. ETHERNET)
17 data from SONET/SDH with provisions for VCAT and LCAS is shown in
18 prior art Figure 1. This is the receive (or sink) portion of the
19 SONET/SDH link. The Front End block 12 performs the well known
20 SONET/SDH functions of descrambling, framing, pointer tracking and
21 alignment. The VCAT & LCAS block 14 extracts VCG differential
22 delay information as well as VCG configuration information
23 embedded in one of the SONET/SDH overhead bytes (H4 in high order,
24 K4 in low order). This VCG information together with the payload

1 is then passed to the Deskew Control block 16. The Deskew Control
2 block 16 performs two functions. First, it calculates the amount
3 of differential delay experienced by each tributary, based on the
4 H4/K4 frame counter. Then it stores each tributary in the Deskew
5 FIFO 18 for a different amount of time based its differential
6 delay. Tributaries that are experiencing short network delays
7 must be stored longer to compensate for those tributaries that are
8 experiencing long network delays. The Demapper block 20 reads the
9 Deskew FIFO 18. Since group members are stored for varying
10 amounts of time in the Deskew FIFO 18, when they are read out,
11 their original phase relationship will be restored. The Demapper
12 block 20 may first need to reorder the group members. Then it can
13 demap the original packet data from each group payload. The
14 output of the Demapper block 20 is packet data destined for the
15 ETHERNET client. Usually, a packet-based elastic storage buffer
16 (Packet FIFO) 22 is provided between the Demapper block 20 and the
17 client interface 24. This Packet FIFO 22 is used to hold packet
18 data during times that the output client interface 24 is blocked,
19 may be used to smooth out bursty reads from the Deskew FIFO 18,
20 and may also be used to avoid running out of client data during a
21 packet transfer.

22

23 Flow control is feasible (although not necessarily
24 recommended) at several points in the system shown in Figure 1.

1 For example, if the local client is temporarily not able to accept
2 data, it can temporarily shut off the output of the client
3 interface 24. This will cause more data to be stored in the
4 Packet FIFO 22. If the Packet FIFO 22 becomes full, it is
5 theoretically possible to shut off the demapping process 20 until
6 more space becomes available in the FIFO 22. This will have the
7 effect of causing data to build up in the Deskew FIFO 18. While
8 this may seem like it takes good advantage of available storage
9 space, it has the drawback of creating head-of-line blocking.
10 Since the Deskew FIFO 18 stores SONET/SDH data with no knowledge
11 of packet boundaries, backpressure at a single client interface
12 port would stop the entire Deskew FIFO output, thus stopping data
13 for all client interface ports, including those that are not
14 blocked.

15
16 It will be appreciated, however, that if the demapping
17 process were stopped and data were allowed to build up in the
18 Deskew FIFO 18, the next logical step would be to shut off data at
19 its input. This, of course, is not possible since the input to
20 the Deskew FIFO 18 is the arriving SONET payloads, which cannot be
21 stopped or slowed down, and there is no mechanism for applying
22 flow control over the SONET network. Therefore, the Packet FIFO
23 22, if present, can be used effectively to prevent data loss due
24 to temporary blockage at the client interface 24 output, but

1 backpressure from the Packet FIFO 22 should not be used to stop
2 the demapping process 20. Since the arriving SONET data cannot be
3 flow controlled, a full Packet FIFO 22 must result in data loss.
4 A well designed system should attempt to minimize this condition.
5

6 The foregoing description on flow control in the sink system
7 assumes that data is moving through the system at a constant rate.
8 Thus, the only need for flow control comes from rate limitations
9 at the output of the client interface 24. When VCG and LCAS are
10 implemented, other flow control problems are introduced. For
11 example, given a static preconfigured VCG that is not changing in
12 size and is composed of three members, it is possible that the
13 first two members arrive with very little network delay, and the
14 third member has a great deal of network delay. In this example,
15 the differential delay between the first two members is zero, but
16 the differential delay between those members and the third member
17 is T . (In practice, T can be as large as 256 msec and still be
18 compensated.) To compensate for the differential delay, the first
19 two members of the VCG need to be stored in the Deskew FIFO 18 for
20 T seconds, while waiting for the third member which is stored
21 minimally while the VCG is demapped. Under these static
22 conditions, the operation of the Deskew FIFO 18 is such that the
23 long term average input and output rates are exactly the same.
24 Instantaneous rates may vary; both the input and output of the

1 Deskew FIFO 18 might be respectively stored and retrieved in high
2 speed bursts followed by long gaps. However, on average, the
3 input rate is exactly the SONET arrival rate. On the output side,
4 data is read from the Deskew FIFO 18 only when it is available.
5 With such an arrangement, the output rate of the Deskew FIFO 18 is
6 paced by the input (arrival) rate. If, however, the third member
7 of the VCG is removed through an LCAS command, the first two
8 members have T seconds worth of data stored in the Deskew FIFO 18
9 now ready to read at the maximum rate the demapping process can
10 support. That rate is possibly greater than the output rate of
11 the client interface 24, so the Packet FIFO 22 begins to fill up.
12 The rapid arrival of T seconds worth of data from several group
13 members can overflow the Packet FIFO, resulting in unnecessary
14 data loss.

15
16 There are several ways to avoid Packet FIFO overflow, but
17 they all have drawbacks. One solution is to always pace the
18 output with the input, regardless of data availability in the
19 Deskew FIFO 18. With such a system in the example given above,
20 after the slowest member is removed from the VCG, data continues
21 to be read out of the Deskew FIFO 18 at the same rate, and T
22 seconds worth of data remains in the Deskew FIFO 18 indefinitely.
23 This results in a system that can only add and never remove delay
24 through the Deskew FIFO 18. Eventually the Deskew FIFO 18 is

1 likely to be filled to its capacity. Operation can continue error
2 free, but the delay through the Deskew FIFO is unnecessary.

3
4 Another solution also paces the output with the input, but
5 when enough data is available in the Deskew FIFO 18, extra reads
6 are performed. This is sometimes done in a fashion similar to
7 SONET pointer movements, thus "leaking" extra data out of the
8 Deskew FIFO 18 at a slow rate. This solution does not guarantee
9 that the Packet FIFO 22 does not overflow as the "leak" times may
10 coincide with a blocked client interface. Also, response times
11 are typically slow, so extra delay through the Deskew FIFO 18 is
12 experienced long after the slow member of a VCG has been removed.

13
14 Still another solution is to size the Packet FIFO 22 large
15 enough to handle all of the possible data from the Deskew FIFO 18.
16 While this solution will not lose any data unnecessarily, it is
17 expensive because the storage requirements for the Packet FIFO 22
18 are typically much less than for the Deskew FIFO 18. This
19 solution also has the effect of temporarily adding undesirable
20 delay through the Packet FIFO 22 while the Deskew FIFO 18 is being
21 emptied.

22
23 The most elegant solutions implement a form of Xon/Xoff flow
24 control between the Packet FIFO 22 and the Deskew FIFO 18. In

1 these systems, depth measurements in the Packet FIFO 22 are used
2 to turn on and off the output of the Deskew FIFO 18.
3 Unfortunately, turning off the Deskew FIFO 18 results in the head-
4 of-line blocking described above, so depth measurements in the
5 Deskew FIFO 18 are sometimes used to turn the flow back on. The
6 control mechanism for these systems is very complex, especially
7 when dealing with multiple client ports.

8 9 SUMMARY OF THE INVENTION

10
11 It is therefore an object of the invention to provide methods
12 and apparatus for preventing or minimizing data loss when sinking
13 a VCG.

14
15 It is also an object of the invention to provide methods and
16 apparatus for preventing or minimizing data loss when removing a
17 member from a VCG.

18
19 It is another object of the invention to provide such methods
20 and apparatus in a manner which avoids the disadvantages of the
21 prior art.

22
23 In accord with these objects which will be discussed in
24 detail below, the methods of the present invention utilize a rate

1 based flow control between the Packet FIFO and the Deskew FIFO and
2 only two rates are used for reading the Deskew FIFO: a maximum
3 rate and an input-limited rate. The maximum rate is the maximum
4 rate of the Demapper and the minimum rate is SONET/SDH arrival
5 rate. A depth measurement of the Packet FIFO is used to determine
6 an upper threshold that is reached when the FIFO is nearly full.
7 Whenever the Packet FIFO depth is below the upper threshold, the
8 demapping process is allowed to run at maximum rate. Once the
9 upper threshold is reached, the Deskew FIFO continues to operate,
10 but the output rate is limited to the input rate. While the upper
11 threshold is reached, it is acceptable for short output bursts to
12 occur faster than the input rate, but the average output rate must
13 be no faster than the input rate. If high speed output bursts are
14 permitted, it is desirable to set the upper threshold of the
15 Packet FIFO at least one burst below full. Efficient
16 implementations should keep the burst size small as well.
17 Although hardware implementations of this method may vary, the
18 method is simple enough that it may be carried out entirely in the
19 Demapper block with fullness information from the Packet FIFO and
20 timing information from the network input. When implementing the
21 methods of the invention, if the client interface output is
22 blocked, the Packet FIFO will continue to fill up and will
23 eventually overflow, but this is an acceptable loss of data, not
24 related to SONET/SDH, VCG, or LCAS processing. No backpressure is

1 ever applied to the Deskew FIFO, so head-of-line blocking is
2 completely eliminated. The entire Packet FIFO is used to its full
3 advantage, and once it is full, if its output is still blocked,
4 the proper response is to drop packets without stopping the
5 demapping process.

6
7 Additional objects and advantages of the invention will
8 become apparent to those skilled in the art upon reference to the
9 detailed description taken in conjunction with the provided
10 figures.

11
12 BRIEF DESCRIPTION OF THE DRAWINGS

13
14 Figure 1 is a simplified schematic diagram of a generic prior
15 art arrangement for recovering packetized data from a SONET/SDH
16 stream;

17
18 Figure 2 is a view similar to Figure 1, illustrating a
19 generic hardware implementation of the methods of the invention;
20 and

21
22 Figure 3 is a more detailed diagram illustrating the
23 presently preferred embodiment of the invention.

1 DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

2
3 Referring now to Figure 2, a system 110 for recovering
4 packetized (e.g. ETHERNET) data from SONET/SDH with provisions for
5 VCAT and LCAS according to the invention is shown. The device 110
6 is similar in some respects to the device 10 described above and
7 similar reference numerals (increased by 100) refer to similar
8 functional blocks. The Front End block 112 performs the well
9 known SONET/SDH functions of descrambling, framing, pointer
10 tracking and alignment. The VCAT & LCAS block 114 extracts VCG
11 differential delay information as well as VCG configuration
12 information embedded in one of the SONET/SDH overhead bytes (H4 in
13 high order, K4 in low order). This VCG information together with
14 the payload is then passed to the Deskew Control block 116. The
15 Deskew Control block 116 performs two functions. First, based on
16 the H4/K4 frame counter, it calculates the amount of differential
17 delay experienced by each tributary. Then, it stores each
18 tributary in the Deskew FIFO 118 for a different amount of time
19 based on its differential delay. Tributaries that are
20 experiencing short network delays must be stored longer to
21 compensate for those tributaries that are experiencing long
22 network delays. The Demapper block 120 reads the Deskew FIFO 118.
23 Since group members are stored for varying amounts in the Deskew
24 FIFO 118, when they are read out, their original phase

1 relationship will be restored. The Demapper block 120 may need to
2 reorder the group members, and then can demap the original packet
3 data from each group payload. The output of the Demapper block
4 120 is packet data destined for the ETHERNET client. Usually, a
5 packet-based elastic storage buffer 122 is provided between the
6 Demapper block 120 and the client interface 124. This Packet FIFO
7 122 is used to hold packet data during times that the output
8 client interface 124 is blocked, may be used to smooth out bursty
9 reads from the Deskew FIFO 118, and may also be used to avoid
10 running out of client data during a packet transfer.

11
12 According to the methods of the invention, the operating rate
13 of the Demapper 120 is adjusted by the Rate Adjustment block 128.
14 The Rate Adjustment block receives a fullness measure from the
15 Fullness Measure block 126 which measures the depth of data in the
16 Packet FIFO 122, and a timing signal 130 from the SONET/SDH Front
17 End 112. When the fullness of the buffer 122 is below a selected
18 threshold (based on the size of the buffer and the characteristics
19 of the SONET/SDH signal), the Demapper 120 is set to run at a rate
20 faster than the SONET/SDH signal input rate. When the fullness
21 reaches the threshold, the demapper rate is slowed to
22 substantially the same as the SONET/SDH signal input rate.

23

1 Figure 3 is a more detailed diagram of a presently preferred
2 implementation of the invention. As seen in Figure 3, the
3 invention is implemented within the Deskew FIFO 218, the Demapper
4 220, and the Packet FIFO 222. The Deskew FIFO 218 includes a
5 Demultiplexing and Control circuit 218a which feeds separate
6 elastic storage 218b for each tributary. The Demapper 220
7 includes a Tempo circuit 220a which receives a write signal WR
8 from the Demux and Control block 218a and a fullness measure Avail
9 from the Deskew FIFO 218. The Avail signal indicates to the
10 Demapper 220 that data is available for demapping and the Demapper
11 220 uses this signal to enable the Reordering and Multiplexing
12 block 220b to output packets via queues 220c. The output of the
13 Demapper 220 feeds elastic storage 222a in the Packet FIFO 222.
14 Depth measurements for each elastic storage member 222a are fed to
15 a comparator 222b which compares the depth measurements to
16 programmed thresholds 222c. The comparator 222b supplies an
17 Almost Full signal for each elastic storage 222a to the Tempo
18 circuit 220a.

19
20 During normal operation, the Demapper 220 reads data from the
21 Deskew FIFO 218 whenever the Avail signal is asserted. This
22 allows the Demapper to operate at its fastest speed. When
23 congestion occurs in the Packet FIFO 222, the Almost Full signal
24 is asserted. This causes the Tempo Circuit 220a to operate the

1 Demapper only when the WR signal is asserted, i.e. at the same
2 speed at which data is written to the Deskew FIFO 218, the
3 SONET/SDH input speed. This assures that no packets are lost in
4 the Packet FIFO due to SONET/SDH signal processing. It will be
5 appreciated, however, that the Client Interface 224 may refuse
6 packets and cause the Packet FIFO to overflow and lose packets.
7 This is the very nature of the ETHERNET protocol that sometimes
8 packets get lost and must be retransmitted.

9
10 There have been described and illustrated herein several
11 embodiments of methods and apparatus for extracting packetized
12 data from a SONET/SDH signal. While particular embodiments of the
13 invention have been described, it is not intended that the
14 invention be limited thereto, as it is intended that the invention
15 be as broad in scope as the art will allow and that the
16 specification be read likewise. Thus, while certain blocks of the
17 invention were described as hardware blocks, it will be
18 appreciated that the various portions of the invention may be
19 implemented in hardware, software, firmware, or a combination of
20 any of the three. In addition, while the invention was described
21 as utilizing FIFOs, it will be appreciated that such FIFOs can be
22 implemented in any of a number of manners, and that other memory
23 elements such as RAMs may be utilized. It will therefore be
24 appreciated by those skilled in the art that yet other

- 1 modifications could be made to the provided invention without
- 2 deviating from its spirit and scope as so claimed.